Learning About Future Oil Prices

Sylvain Leduc — *FRBSF*
Kevin Moran — *Université Laval*
Robert Vigfusson[*†] — *Federal Reserve Board*

March 27, 2013

1 Introduction

Between 1990 and 2002, the implicit forecast of the futures market was that the long-run price of oil would be near $20 per barrel. The futures market failed to predict that the developments at the start of the 21st century would radically change the outlook for oil prices. Whereas many commentators have attributed the associated forecasts errors to speculation or market inefficiency, this paper provides an explanation of the movements in oil prices based on learning.

In the early 2000s, concerns about geopolitical unrest (especially in the Middle East), strong income growth (especially in China and other emerging markets), and modest oil production growth led to sharp upward movements in spot oil prices. The price of WTI increased from $20 to $30 per barrel over the course of 2002. Initially, investors were surprised by the persistence of the oil-price increases. Although spot prices soared, the futures market price of oil delivered two years hence rose only modestly. The implicit forecast was that the upward pressure on oil prices would only be transient. Moreover, many official agencies, such as the Energy Information Agency in the United States, shared the view that the run-up in oil prices would be temporary.

Although the increases in oil prices were first perceived as temporary, this outlook did not last. Following a decade of continued predictions that oil prices would revert back to a long-term trend, far-dated futures prices began to move one for one with the spot price. By the middle of 2008, the futures curve was flat, with the spot price and the far-dated price nearly equal and moving together.

Subsequently, both spot and futures oil prices have remained elevated but the tight link between the spot price and the far-dated futures price has relaxed somewhat. For

[*] The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, any other person associated with the Federal Reserve System.

[†] Comments and suggestions can be directed to robert.j.vigfusson@frb.gov.

instance, the Great Recession drove down the spot price of oil much more than the farther-dated futures. Likewise, the futures curve for Brent oil prices been sharply downward sloping, with implicit predictions of over 10 percent decline over the next two years.

In this paper, we show that the developments in the oil futures markets over the past 20 years are consistent with investors learning about whether underlying shocks are transitory or permanent. We provide evidence using the Kalman filter to infer the permanent and transitory components of shocks to spot oil prices, we show that this form of learning can then explain remarkably well the observed behavior of futures prices. Our simple framework accounts for the relatively slow increase in oil-price futures at the beginning of the past decade and their unprecedented run-up between 2004 and 2008. Even during the first half of 2008, a period during which oil prices reached historic highs, the model predicts a level of futures prices that is in line with the data.

We also contrast our baseline result to that obtained under a constant-gain learning process under which recent developments are weighted more heavily than earlier data. We show that the fit of the model is significantly worse under constant-gain learning, predicting much larger upward movements in futures prices during the 2003-2005 period than those actually observed. Thus, our results indicate that participants in the oil markets continued to place substantial weight on developments in the 1980s and early 1990s when forming their expectations of future oil prices.

We then examine the implication of learning for the efficiency of oil futures markets, by estimating simple forecasting-efficiency regressions over different sample periods. Although there is little indication of systematic forecast errors over the period from 1990 to 2012, our results indicate that futures prices were often a biased predictor of future spot prices that during the previous decade. We show that the kind of mistakes made over the past 10 years or so are consistent with market participants being surprised by the persistence of movements in the spot price of oil and adjusting their prior of the relative importance of temporary and permanent shocks.

Our work complements a growing literature interested in understanding these large movements in the spot price of oil and its changing relationship with the futures market. In particular, it relates to recent work on the role of financial markets in driving oil prices.[1] For instance, Hamilton and Wu (2012)argue that increased participation by index-fund investing has reduced the oil futures premia since 2005, accounting for the smaller gap between spot and futures prices observed in the data between 2005 and 2008. Similarly, recent work by Buyuksahin, Haigh, Harris, Overdahl, and Robe (2008)argue that increased market activity by commodity swap dealers, hedge funds, and other financial traders, has helped link crude oil futures prices at different maturities. As such, these papers attribute the recent changes in the futures markets to the increased financialization of commodity markets, while we show that these movements can be largely attributed to learning.

[1] This literature is large and growing. Some of the many papers discussing this issue include Singleton (2012), Irwin and Scott (2012) and Fattouh, Kilian and Mahadeva (2012), Kilian and Murphy, Alquist and Gervais (2012).

The rest of the paper is organized as follows. After describing in more detail developments in oil prices, we then layout our theoretical model. We then illustrate how uncertainly about the persistence of a shocks alters the model's responses relative to the same model but where the agents know with certain the persistence of the shock. Given this theoretical background, we then estimate a time series model of permanent and temporary shocks to oil prices. Given our assumptions regarding learning and using only spot price data, we estimate that the predicted permanent component of spot prices matches very well the observed developments in futures markets. We then show that the empirical evidence that has been interpreted as evidence of a time-varying risk premium is also consistent with our learning-based explanation.

2 Oil Prices Over the Past Two Decades

In this section, we present evidence suggesting that participants in the oil market may have been surprised by the degree of persistence of underlying shocks hitting the world economy. We first examine the behavior of spot and futures oil prices over the past two decades, before looking at forecasts from the U.S. Department of Energy.

Consider the movements in the spot and futures prices of oil since the early 1990s that are depicted in Figure 1. During the last decade of the previous century, the spot price of oil tended to gyrate around fairly steady oil-price futures, suggesting that market participants view economic developments impacting the oil market as mainly temporary. Underlying shocks would tend to move the spot price of oil, at times substantially, but the spot price would tend to rapidly return to roughly $18 per barrel, the level of futures prices for most of that period. Clearly, whatever the disturbances affecting the world economy, market participants did not view them as persistent enough to alter their view of oil prices in the future.

As shown in Figure 2, during the 2000-2005 period, the relationship between spot and futures price changed compared to the 1990s. For most of that period, the spot price of oil rose and remained well above futures prices. Oil-price futures remained initially low, consistently fluctuating below $20 per barrel until 2003, before initiating a steady rise that brought futures prices to roughly $50 per barrel by the mid-2000s. To us, this pattern suggests that market participants initially thought that movements in spot prices would likely be temporary, as indeed was the case throughout the 1990s. To be sure, the large increases in the spot price of oil were matched by much more muted increases in futures prices, between 2000 and 2003.

This clearly changed in 2004 when futures prices started a rapid and sustained rise that converged to the level of spot prices within a year. Spot and futures prices then tended to move in lockstep between 2005 and 2008. Our interpretation of these movements is that after being consistently surprised by the persistence of the increases in spot prices, market participants reassessed their views about the persistence of underlying shocks, placing more weight towards more persistent or permanent developments.

The financial crisis of 2008-2009 and the following deep economic contraction around the world provide an interesting event to examine the plausibility of our interpretation. The collapse in economic activity in the fall of 2008 sent the spot price of oil tumbling, losing more than 50 percent of its value. However, the market perceived the extent of the fall as likely temporary and, in contrast, oil-price futures declined much more moderately. Since then, spot prices have risen to the level of futures prices, which has been relatively constant at slightly more than $90 per barrel since 2010.

Figure 2 shows that our view of the developments in the oil market over the past two decades is consistent with the paths of futures prices. This figure presents the expected path of futures prices as of particular dates since 1990. It shows that for dates between 1990 and 2003, the paths were mostly downward sloping, with the slopes being more pronounced during the 1990s. Moreover, these paths clearly indicate that the market expected futures prices to remain reasonably constant in the medium run at around $18 per barrel.

As shown in Figure 2, this pattern dramatically changed around 2003. From then on, the paths of oil price futures tend to be upward sloping, consistent with the view that market participants view economic developments as having more persistent effects on oil prices. We once again briefly observe downward sloping futures paths during 2008 and also more recently.

3 Empirical framework

We develop a simple unobserved components model to illustrate the role of permanent and temporary shocks in determining oil prices. To illustrate how the futures markets may have learned about the relative importance of permanent and temporary shocks, we present a simple time series model for oil prices. Although this model does not include the cross-equation linkages that are implied by the DSGE model, this simpler model will illustrate the paper's main idea about learning without having to impose any counter-factual cross equation restrictions.

Suppose that the price of oil p_t (expressed in logs) is the following linear combination of a permanent component e_t^P and a temporary component e_t^τ.

$$p_t = e_t^P + e_t^\tau$$

The permanent component is modeled as a random walk with drift.

$$e_t^P = \mu + e_{t-1}^P + v_t$$

where v_t is i.i.d. normal with variance σ_p^2. The temporary component, is assumed to be an AR(1) process

$$e_t^\tau = \phi_\tau e_{t-1}^\tau + \varepsilon_t$$

where ε_t is i.i.d normal with variance σ_τ^2 and $|\phi_\tau| < 1$. This model corresponds to an unobserved components model, which has been applied frequently in macroeconomics.

For a review, see Perron and Wada (2009). In addition Schwartz and Smith (2000) applied this model to oil prices in the 1990s. Unlike the current paper, they did not consider the possibility of structural change, which is not surprising given their data sample.

Under the assumption that the model is quarterly and given full information on the temporary and permanent components, the equations for one-year ahead futures price $f_{t,4}$ and two-year ahead futures price $f_{t,8}$ are

$$\begin{aligned} f_{t,4} &= E_t p_{t+4} = 4\mu + e_t^P + \phi_\tau^4 e_t^\tau \\ f_{t,8} &= E_t p_{t+8} = 8\mu + e_t^P + \phi_\tau^8 e_t^\tau . \end{aligned}$$

Absent full information on the current levels of e_t^P and e_t^τ, these futures will instead be based on the best forecasts given the observed behavior of p_t.

$$\begin{aligned} f_{t,4} &= E_t p_{t+4} = E_t \left(4\mu + e_t^P + \phi_\tau^4 e_t^\tau | \{p_{t-i}\}_{i=t}^1\right) \\ f_{t,8} &= E_t p_{t+8} = E_t \left(8\mu + e_t^P + \phi_\tau^8 e_t^\tau | \{p_{t-i}\}_{i=t}^1\right) \end{aligned}$$

With a small value of μ and moderate values of ϕ_τ, this model can imply the downward sloping futures curve arising from a positive temporary shock. These futures can be calculated using the Kalman filter, which we now discuss.

3.1 Discussion of the Kalman Filter

This model can be easily written in a form appropriate for the Kalman filter. The model consists of the following parameters: the standard deviation of the shock to the temporary component σ_τ , and the standard deviation of the shock to the permanent component σ_p and the persistence parameter of the temporary component ϕ_τ. The Kalman filtering equation relates how the observed variables y_t (such as prices) respond to the changes in the unobserved state vector ξ_t which is comprised of the trend and the temporary and permanent components.

$$\xi_t = \begin{pmatrix} \mu \\ e_t^P \\ e_t^\tau \\ \mu \\ e_{t-1}^P \\ e_{t-1}^\tau \end{pmatrix} .$$

The equations for the dynamics for the observed variables y_t are given by the following system

$$\begin{aligned} \xi_t &= F\xi_{t-1} + \varepsilon_t \\ y_t &= H\xi_t \end{aligned}$$

Given an initial estimate of $\xi_{t|t-1}$ and $P_{t|t-1}$, [the mean squared error matrix equation 13.2.6 in Hamilton (1994)] we can use Hamilton's updating equation [13.2.23 page 381]

$$\xi_{t|t} = \xi_{t|t-1} + P_{t|t-1} H (H' P_{t|t-1} H)^{-1} \left(p_t - H' \xi_{t|t-1}\right) \tag{1}$$

The news is the value of $P_{t|t-1}H(H'P_{t|t-1}H)^{-1}\left(y_t - H'\xi_{t|t-1}\right)$. One uses the forecast error$\left(y_t - H'\xi_{t|t-1}\right)$ to update their estimates of the size of the permanent and transitory component. The value of $P_{t|t-1}$ governs the extent to which a given surprise is assumed to be part of the permanent component relative to the transitory component. The values of σ_v^2 and σ_ε^2 play a large role in determining whether an observed spot price increase is interpreted as permanent. If a one percent increase in spot price increase is interpreted as transitory, then $E_t\left(p_{t+k}|\left\{p_{t-i}\right\}_{i=t}^1\right)$ only increases by ϕ_τ^k. In contrast, if the same increase is interpreted as permanent, then the value of $E_t\left(p_{t+k}|\left\{p_{t-i}\right\}_{i=t}^1\right)$ increases by the same amount. If a shock is actually permanent but mistaken for temporary, then the value of $E_t\left(p_{t+k}|\left\{p_{t-i}\right\}_{i=t}^1\right)$ (and hence the futures price) will include this error.

3.2 Estimation of the Model Parameters.

If the model parameters were known with certainty than learning about $E_t\left(p_t|\left\{p_{t-i}\right\}_{i=t}^1\right)$ would be a simple application of Equation 1. However, if instead we assume that the model's structure is known but not its parameters values, then these parameter values must be estimated given lagged observed data. $\left\{p_{t-i}\right\}_{i=t}^1$

The standard approach would be to assume that the model parameters are constant and therefore are estimated using the standard likelihood function.

$$LL_T = -\sum_{t=1}^{T}\left(\frac{1}{2}\ln 2\pi + 0.5\log\|V_t\| + (p_t - Ep_t)\,V_t^{-1}\,(p_t - Ep_t)\right) \tag{2}$$

where $V_t = H'P_{t|t-1}H$ the variance of the prediction errors.

An alternative approach is to suppose that estimation is concerned with time-variation in the model parameters. The estimation procedure puts more weight on recent observations. In the learning literature, this approach is referred to as constant gain learning. For an example, see the recursive least squares algorithm in Cho, Williams and Sargent (2002). In the spirit of this work, we modify the likelihood function as follows.

$$LL_T = (1-\chi_T)\,LL_{T-1} - \chi_T\left(\frac{1}{2}\ln 2\pi + 0.5\log\|F_t\| + (p_t - Ep_t)\,V_t^{-1}\,(p_t - Ep_t)\right) \tag{3}$$

If $\chi_t = \frac{1}{t}$ then all observations have the same weight, equivalent to the standard likelihood function described above. In contrast if χ_t is a constant, then recent observations are more important than lagged observations. In particular, for a dataset of T observations, the first observation contributes

$$\prod_{t=1}^{T}(1-\chi_t)\,\chi_1$$

whereas the most recent observation (observed at time T) has a much greater weight of

$$\chi_T.$$

3.3 Our Three Step Procedure

In order to estimate the value of

$$E_t\left(8\mu + e_t^P + \phi_\tau^8 e_t^\tau | \left\{p_{t-i}\right\}_{i=t}^1, \Gamma_{t-1}\right)$$

we apply the following three-step procedure. In the first step, we use prices observed up to time t-1 (the previous quarter) and estimate model parameters Γ_{t-1} using the standard likelihood function.

The second step is to apply the Kalman filter using these estimated model parameters and observed prices through time t (ie the current quarter) to get e_t^p and e_t^τ. The third step uses these estimated e_t^p and e_t^τ and Γ, to construct f_t, k

Our three step procedure is likely not optimal. Instead, we should use nonlinear methods (such as a particle filter) to address the problem of learning about Γ and e_t^p and e_t^τ jointly. However, our procedure has the benefits of simplicity and clarity.

3.3.1 Variance Contribution.

As a simple statistic to measure the importance of permanent shocks, we note that the equation for $\triangle p_t$ is the following

$$\begin{aligned}\triangle p_t &= \left(e_t^P + e_t^\tau\right) - \left(e_{t-1}^P + e_{t-1}^\tau\right) \\ &= \mu + \varepsilon_t^P + \left(\phi_\tau - 1\right) e_{t-1}^t + \varepsilon_t^t\end{aligned}$$

Therefore, we have that the variance of $\triangle p_t$ can be expressed as

$$var\left(\triangle p_t\right) = \sigma_p^2 + \sigma_\tau^2 + \frac{\left(1-\phi_\tau\right)^2}{1-\phi_\tau}\sigma_\tau^2.$$

As such the fraction of the growth rates' variance that is due to permanent shocks can be calculated as

$$\frac{\sigma_p^2}{\sigma_p^2 + \sigma_\tau^2 + \frac{(1-\phi_\tau)^2}{1-\phi_\tau}\sigma_\tau^2}.$$

4 Empirical Results

In this section, we will estimate our model using the average monthly price at the end of each quarter for the spot price of WTI oil. We will reserve the futures prices to test our specification. In particular, we will study whether the model-implied futures estimates Ep_{t+8} match the two-year ahead futures price observed on the business day closest to one week following the end of the quarter. We use this timing so that our price forecast would be consistent with the information available when the actual futures price was observed.

As described above, the first step is to estimate Γ_t where

$$\Gamma_t = \left\{ \begin{array}{cccc} \mu & \phi & \sigma_p^2 & \sigma_\tau^2 \end{array} \right\}.$$

Figure 3 reports the estimated value of Γ_t that maximize the standard log likelihood function using a recursively constructed sample. The solid black lines report the estimated coefficients. The grey intervals are 2 standard-deviation confidence intervals. The results are in line with the narratives for Figures 1 and 2. The top left panel of Figure 3 reports the value of μ. Using just the pre-2000 part of the sample, the value of μ is slightly below zero, which implies a negative trend for the nominal price. However, as the estimation sample includes more of the twenty-first century, the estimated trend first turns positive and then begins to increase. The maximum value of the trend is 0.018, which is estimated for a sample that ends in the second quarter of 2008. Along with spot oil prices, the estimated trend then declines and stabilizes around 0.013 or at an annual rate of just over 5 percent per year. In contrast to μ, the value of ϕ (the temporary component's AR parameter) is relatively much more stable. The value of ϕ is first estimated at 0.63. Its estimated value peaks at 0.78 and then declines to 0.68. Even for a large value of ϕ_τ, the value of $E_t p_{t+8}$ is not strongly affected by a change in the temporary component. Since ϕ_τ is less than one, it decays quite quickly. In particular, a one percent increase in the temporary component results in the value of $E_t p_{t+8}$ increasing by only 0.13 (ϕ_τ^8) percent. The bottom panels of Figure 3 report the estimation results for σ_p and σ_τ. These estimated coefficients do vary as the sample included in the MLE estimation changes. In particular, the estimated value of σ_p is zero for the initial sample. However, for the sample ending in the second quarter of 2008, the estimated value of σ_p peaks at 0.118, and then declines as additional observations are added to the sample, with the final value being 0.089. The estimated value of σ_τ follows the opposite pattern. It begins high, starting at 0.163 and then declines to 0.107 by 2008 and then increases back to 0.149.

The evolution of these model parameters can also be viewed by considering the role of permanent shocks in affecting the variance of $\triangle p_t$, which is reported in Figure 4. As indicated by the black line, the estimated contribution of permanent shocks only slowly increases over time. In the early part of the sample, because the standard deviation of innovations to the permanent component is extremely small, the permanent component's contribution to the var($\triangle p_t$) is negligible. Therefore, although the temporary component does not strongly affect the futures price, in the end-of-2001 sample, the temporary component is the main driver of changes in the spot price. These results are very much in line with Figure 1, where the futures curves show transitory increases and prices returning to a long-term trend. The estimated contribution of permanent shocks actually peaks at 40 percent in the second quarter of 2008 along with spot oil prices. The sharp fall in oil prices in the last half of 2008 resulted in a somewhat lower estimates of the role of permanent shocks.

Our next exercise is to determine what would be the values of the two-year (eight quarter) futures prices $f_{t,8}$ that a contemporaneous econometrician would estimate using the Kalman filtering formula Equation 1 with the estimated value of Γ_{t-1}.

Figure 5 illustrates the evolution of the futures prices. Comparing the actual futures prices (the thick black line) with the estimated value of $E_t\left(p_{t+8}|\left\{p_{t-i}\right\}_{i=t}^{1},\Gamma_{t-1}\right)$ (the purple dots) is highly instructive. The estimated values of $E_t p_{t+8}$ matches very well the actual futures prices. Most of the time the futures price and the estimated value of $E_t\left(p_{t+8}|\left\{p_{t-i}\right\}_{i=t}^{1},\Gamma_{t-1}\right)$ matches very closely.[2] In the early 2000s, actual futures prices are well below the observed spot prices (the orange line). However, both the actual futures prices and the estimated values of $E_t p_{t+8}$ move together and by mid-2008 are very close to the spot prices. The speed of learning in the futures market could be said to be equivalent or perhaps, given the slight discrepancies in 2005 and 2006, even faster than the learning in our time-series model. Although the learning appeared to be somewhat faster than the MLE estimates, the futures market did not become excessively responsive to recent price increases. Thereafter the spot price declines much more than either the actual futures prices or the value of $E_t\left(p_{t+8}\right)$. At the end of the sample, $E_t\left(p_{t+8}\right)$ is somewhat above the value of the actual futures curve. It is suggestive that the estimated trend $\hat{\mu}$ may be larger than the trend embodied in the actual futures prices.

Figure 6 illustrates the importance of time variation of Γ_{t-1}. This Figure replicates from Figure 6, the time series of spot prices, futures prices, and expected future prices $E_t\left(p_{t+8}|\left\{p_{t-i}\right\}_{i=t}^{1},\Gamma_{t-1}\right)$. In addition, Figure 6 reports the expected futures prices $E_t\left(p_{t+8}|\left\{p_{t-i}\right\}_{i=t}^{1},\Gamma_{2002Q4}\right)$ (the light blue dots) where current spot prices are used in the Kalman filtering exercise; but, where the model coefficients are fixed at the estimated values at the end of 2002 (Γ_{2002Q4}). Using fixed parameter values, the Kalman filter generates a much smaller increase in the expected spot price. In particular, the value of $E_t\left(p_{t+8}|\left\{p_{t-i}\right\}_{i=t}^{1},\Gamma_{2002Q4}\right)$ for t=2008Q2 only reaches \$76 per barrel in 2008 when the spot price is \$135 per barrel. In contrast, given the same data but using the estimated coefficients for a sample that ends through 2008Q1, the value of $E_t\left(p_{t+8}|\left\{p_{t-i}\right\}_{i=t}^{1},\Gamma_{t-1}\right)$ is estimated at nearly \$130 per barrel.

4.1 Constant Gain MLE Results

Given the importance of time variation, our next exercise is to consider whether there exists a value of χ_T such that the weighted maximum likelihood estimation results in a better match of the observed futures prices. As shown in Figure 6, the value of χ_T that best matches the 2-year ahead futures path is 0.015. Using such a weight, an observation 11 years in the past gets only half as much weight in the likelihood as today's observation.

[2] The grey bands indicate the confidence interval which is defined as the following set

$$\left\{E_t\left(p_{t+8}|\left\{p_{t-i}\right\}_{i=t}^{1},\tilde{\Gamma}_{t-1}\right)|\left(\Gamma_{t-1}-\tilde{\Gamma}_{t-1}\right)' W\left(\Gamma_{t-1}\right)^{-1}\left(\Gamma_{t-1}-\tilde{\Gamma}_{t-1}\right)\leq 4.5\right\}$$

where Γ_{t-1} is the mle estimate and $W\left(\Gamma_{t-1}\right)$ is the corresponding estimated variance covariance matrix. The critical value of 4.5 is chosen as the 66 percentile of the chi-squared distribution with 4 degrees of freedom.

In contrast, a larger value of χ_T pushes up the value of $E_t\left(p_{t+8}\right)$ dramatically. When χ_T equals 0.03, then we have that an observation 11 years in the past gets only get 25 percent of the weight of today's observation. Large values of χ_T result in a much higher oil price forecast. Ex post, such a forecast would have been more accurate in the sample through early 2008. Thereafter, the forecast is too aggressive and is too high relative to realized prices. The constant-gain forecast of $180 per barrel in the spring of 2008 would have resulted in some very costly bets, given that prices have remained well below that value.

The importance of χ_T in influencing the model parameters is seen in Figure 7. The Figure repeats the mle estimates from Figure 3. In addition. the Figure reports the estimated coefficients for the cases when χ_T equals 0.015 and when χ_T equals 0.03. The main effect of increasing χ_T is to result in higher estimates of the trend. In addition, as shown in Figure 8, the role of permanent shocks does increase when the model is estimated with a constant gain likelihood function.

To summarize the results so far, we have shown that a very simple learning model can readily match the observed patterns in the futures market. Even though the model is estimated using only spot price data, the model matches well two-year ahead futures prices. This ability to match observed outcomes suggests that the short-comings of the futures markets can be attributed to information processing challenges rather than to market inefficiency or to the influx of speculators generating a risk premium. The next section explores the evidence for a risk premium and highlights that the empirical evidence in favor of a risk premium can just as well be taken as evidence for our model.

4.2 Risk Premium

In contrast to our emphasis on learning, the alternative approach in discussing shortcomings of the futures markets is to emphasis a role for the risk premium. A risk premium would imply that the futures price does not equal the expected spot price in order to compensate market participants for participating in the futures market. The futures price is

$$f_{t,k} = E_t\left(p_{t+k} \middle| \left\{y_{t-i}\right\}_{i=t}^{1}, \Gamma\right) + \Phi_t$$

where Φ_t is the compensation required to balance demand and supply in the futures market.

One way to test for a risk premium is the classic forecast efficiency regression.

$$p_{t+k} - p_t = \alpha + \beta\left(f_{t,k} - p_t\right) + u_{t+k} \tag{4}$$

The estimated value of β can be show to equal (asymptotically)

$$\beta = \frac{var\left(E_t p_{t+k} - p_t\right)}{var\left(E_t p_{t+k} - p_t + \Phi_t\right)} + \frac{cov(E_t p_{t+k} - p_t, \Phi_t)}{var(E_t p_{t+k} - p_t + \Phi_t)} + \frac{cov(u_{t+k}, E_t p_{t+k} - p_t + \Phi_t)}{var(E_t p_{t+k} - p_t + \Phi_t)}$$

or else

$$\beta = 1 - \frac{cov(\Phi_t, f_t - p_t)}{var(f_t - p_t)}$$

and the value of α

$$\alpha = (1-\beta)\,E\left(p_{t+k} - p_t\right) - \beta E\Phi_t$$

Under the null hypothesis of no risk premium. β equals one and α equals zero. Of course, alternative empirical strategies for testing the risk premium exist. In particular, Hamilton and Wu (2012) apply a minimum chi-squared estimation strategy to their structural model to estimate a risk premium. However, as we make clear below, our results are very similar in spirit to those reported in Hamilton and Wu.

4.2.1 Empirical Results

We estimate equation 4 using NYMEX futures data, with the spot price being measured as the front-month futures value on the business day closest to the 15th of each month. The standard errors are constructed to correct for heteroskedasticity and serial correlation. For the full sample, the point estimates are α equals 0.09 and β equals 0.9. As reported in Table 1, over the full sample, one would fail to reject the null hypothesis of no risk premium.

Although the full-sample results do not provide evidence in favor of a risk premium, many papers in the literature have reported time variation in the evidence of the risk premium. In particular, Hamilton and Wu (2012) report that the compensation to the long position is smaller on average; but, more volatile in the recent data since 2005 relative to the earlier period of between 1990 and the end of 2004. Although they find a strong evidence of a break between the two samples, they do not formally determine the date of the breakpoint.

Because Hamilton and Wu's results are for futures contracts with shorter horizons, their results are not directly comparable with our own. However, we also find a large degree of time variation in the estimated coefficients. Figure 9 reports the recursive sample estimation of β. These results are suggestive of time variation. However, even more time variation can be found by breaking up the data into subsamples. We report results for three subsamples: the Early Sample (1990-2002), the Middle Sample (2003-2007), and the Late Sample (2008-2012). These samples are chosen to highlight the degree of variation in the estimated coefficients.

As reported in Table 1, for the Early Sample of 1990-2002, the point estimates of the intercept and slope are even closer to zero and one respectively, which strongly suggests no risk premium before 2002.[3] In contrast, for the Middle Sample (2003-2007), we would

[3] The reported end point in these split samples is the date of the last evaluated forecast; so in the Early Sample 1990-2002, the last value of the dependent variable is the difference between the December 2002 futures contract observed on November 15 2002 and the December 2001 futures contract observed on November 15 2001.

strongly reject the null hypothesis of no risk premium, with an estimate of β of -0.14. Finally, for the Late Sample, the estimated β is 2.89.

The extreme estimates for β for the Middle and Late Sample might seem challenging to explain. One could dismiss them as reflecting just sampling uncertainty and, in fact, a combined sample finds a more reasonable estimate of β equal to 0.63. However, assuming that these values are comparable to their population values, a negative β would require that the $cov(E_t p_{t+k} - p_t, \Phi_t)$ be negative. When equation 4 is estimated with exchange rates rather than oil prices, negative β are very common (see Engel 1996 for a survey). As such, one might try to apply theories developed there to explain these results (for example see Burnside Eichenbaum and Rebelo 2009). In contrast, the greater than one estimates for the Late Sample are not frequently observed in other empirical literatures. Hence explanations for this result due to risk premium would be more challenging. Although we would reject the null hypothesis of β equals to one in the Late sample and hence reject the null hypothesis of no risk premium, an estimate of β of 2.89 would be challenging to reconcile with the standard risk premium story. As we explain below, although these results are challenging to attribute to a risk premium, we can explain these results using our learning model.

Table 1: Estimated Forecast Efficiency Coefficients

	Constant	Slope	T-tests			Wald Test*
	α	β	$\alpha = 0$	$\beta = 0$	$\beta = 1$	$[\alpha = 0, \beta = 1]$
Full Sample (1990-2012)	0.09	0.90	1.88	3.39	-0.40	3.68
Early Sample (1990-2002)	0.07	0.98	1.02	2.92	-0.05	1.54
Middle Sample (2003-2007)	0.19	-0.14	3.16	-0.27	-2.37	43.19
Late Sample (2008-2012)	-0.12	2.89	-0.99	4.49	2.93	10.55
(2003-2012)	0.13	0.63	1.64	1.24	-0.72	2.75

Notes: Regression results for Monthly data, Newey-West Standard Errors, d Test
*Critical Value for Wald Test is 5.99

4.2.2 Reconciling with Learning

Although the results for the late sample would be difficult to reconcile with a risk-premium based explanation, the estimated coefficients for both the Late Sample and the earlier samples can be reconciled with our learning story. First, we note that the equation for the estimated β in the forecast efficiency regression is

$$\beta = 1 + \frac{cov(E_t p_{t+k} - f_t, f_t - p_t)}{var(f_t - p_t)}$$

In a learning story

$$E_t p_{t+k} - f_t = E_t(p_{t+k}|\Gamma) - E_t(p_{t+k}|\hat{\Gamma})$$

is the error you make when you form your expectations using the wrong parameter values. As such, this is how we reconcile the forecast efficiency regression results with the learning story.

As we emphasized above, our learning story has three episodes. In the first, agents correctly estimate the variance of permanent shocks relative to the realized data. In the second, agents underestimate the variance of permanent shocks relative to the realized data. In the third, agents over-estimated the variance of permanent shocks relative to the realized data. As we will now show, this story is consistent with the results reported in Table B.

To illustrate this consistency, we construct three simulation-based scenarios. First, we simulate spot prices using our unobserved component model as the data generating process combined with our full-sample estimates of μ, ϕ, σ_v, and σ_ε. Using the simulated data, we apply the standard Kalman filter updating equation to estimate the implied value of f_t^k. However, we will apply the updating equation in three different scenarios.

The first scenario matches our story about the early sample (1990-2002) that agents during that time period had accurate estimates of the model's parameters. In the first scenario, the futures price is constructed using the dgp's values of μ, ϕ, σ_p, and σ_ε. The second scenario matches our view of the Middle Sample (2003-2007) that market participants had underestimated the variance of the permanent shock. In the second scenario, we use the dgp's values of μ and ϕ. But set $\hat{\sigma}_\varepsilon = \sigma_\varepsilon + 0.04$ and $\hat{\sigma}_p = \sigma_p - 0.04$. In the scenario, agents will perceive increases in prices as more likely due to temporary shocks than to permanent shocks.

In the third scenarios, we match what we think happened during the late sample, where agents overestimated the permanent shock. The agents are assumed to construct futures prices using the dgp's values of μ and ϕ. But set $\hat{\sigma}_\varepsilon = \sigma_\varepsilon - 0.04$ and $\hat{\sigma}_p = \sigma_p + 0.04$

For each scenario calculate

$$f_t^k = E\left(p_{t+k}\right)$$

and then estimate the risk premium regression equation 4. Table 2 reports the estimated values of β for the three different scenarios.

Table 2 Estimated β,under different scenarios.

Scenario	Mean β
True σ_p	1.0
Underestimated σ_p	0.65
Overestimated σ_p	1.92

These simulation-based results are very much in line with our empirical results that are reported in Table 2. When the futures are constructed knowing the true parameters, the estimated value of β is close to one. When the standard deviation of permanent shock is underestimated, then shocks are more likely to be interpreted as temporary shocks. To get a good estimate of the predicted price change, one would need to underweight $f_t - p_t$ and hence the estimated value of β is well below one. Likewise when σ_p is overestimated, all shocks are interpreted as permanent shocks. To get a good estimate of the predicted price change, one would need to overweight $f_t - p_t$ and hence the estimated value of β is nearly two. Although the reported results in Table 2 did not achieve the extreme values reported in Table 1, even greater divergences between the true σ_p and $\hat{\sigma}_p$ will result in more extreme values of β.

5 Conclusions

The futures market failed to predict that the developments at the start of the 21st century would radically change the outlook for oil prices. Whereas many commentators have attributed the associated forecasts errors to speculation or market inefficiency, this paper provides an empirical and a theoretical explanation of the movements in oil prices based on learning. Using this approach, we can match well the observed pattern of the futures market.

References

Alquist, Ron and Lutz Kilian, 2010, "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics* 25(4).

Alquist, Ron, and Olivier Gervais (2011) "The Role of Financial Speculation in Driving the Price of Crude Oil" Working paper, Bank of Canada.

Alquist Ron, Lutz Kilian and Robert J. Vigfusson, 2011, "Forecasting the Price of Oil," International Finance Discussion Paper No. 1022, Board of Governors of the Federal Reserve System.

Andolfatto, David and Scott Hendry and Kevin Moran (2008), "Are inflation expectations rational?" *Journal of Monetary Economics* 55(2): 406–422.

M Bodenstein and L Guerrieri, 2011. *Oil efficiency, demand, and prices: a tale of ups and downs,*

Burnside, Craig, Martin Eichenbaum and Sergio Rebelo. 2009 "Understanding the Forward Premium Puzzle: A Microstructure Approach," *American Economic Journal: Macroeconomics* 1(2), 127-54.

Buyuksahin, Haigh, Harris, Overdahl, and Robe 2008 *Fundamentals, Trader Activity and Derivative Pricing,* Chernenko Sergey V., Krista B. Schwarz and Jonathan H. Wright 2004, “The Information Content of Forward and Futures Prices: Market Expectations and the Price of Risk,” International Finance Discussion Paper No. 808, Board of Governors of the Federal Reserve System.

Chinn, Menzie D., Olivier Coibion, 2010, “The Predictive Content of Energy Futures,”NBER Working Paper No. 15830 (January).

Cho, I.-K., Williams, N., Sargent, T.J., 2002. Escaping Nash inflation. *Review of Economic Studies* 69, 1–40.

Erceg, Christopher J. & Levin, Andrew T., 2003. "Imperfect credibility and inflation persistence," *Journal of Monetary Economics*, Elsevier, vol. 50(4), pages 915-944, May.

Engel, Charles, 1996. "The forward discount anomaly and the risk premium: A survey of recent evidence," *Journal of Empirical Finance*, Elsevier, vol. 3(2), pages 123-192, June.

Etula, Erkko (2010) Broker-Dealer Risk Appetite and Commodity Returns" Working paper, Federal Reserve Bank of New York.

Fattouh, Bassam, Lutz Kilian, and Lavan Mahadeva (2012) “The Role of Speculation in Oil Markets: What Have We Learned So Far?" Working paper, University of Michigan.

J Hamilton 2009 *Causes and Consequences of the Oil Shock of 2007-08*, Brookings Paper

Hamilton, James D. (1994) Time Series Analysis Princeton University Press, Princeton, NewJersey.

Hamilton, James D. and Jing Cynthia Wu 2012 “Risk Premia in Crude Oil Futures Prices” *Manuscript*

Irwin, Scott H., and Dwight R. Sanders (2012) “Testing the Masters Hypothesis in Commodity Futures Markets" Energy Economics 34, 256-269.

Kilian, Lutz, and Daniel P. Murphy (2011) “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil" Working paper, University of Michigan.

King, Robert and Mark Watson (2002), “System Reduction and Solution Algorithms for Singular Linear Difference Systems under Rational Expectations” *Computational Economics* 20: 57–86.

Mou, Yiqun (2010) “Limits to Arbitrage and Commodity Index Investment: Front-Running the Goldman Roll" Working paper, Columbia University.

Pierre Perron, Tatsuma Wada, “Let’s take a break: Trends and cycles in US real GDP", Journal of Monetary Economics, Volume 56, Issue 6, September 2009, Pages 749-765, Pindyck, Robert S., 2001, "The Dynamics of Commodity Spot and Futures Markets: A Primer," *Energy Journal* 22(3): 1-29.

Reichsfeld, David A and Shaun K. Roache, 2011 "Do Commodity Futures Help Forecast Spot Prices?" *IMF Working Paper* No. 11/254:

Schwartz, Eduardo S., and James E. Smith (2000) "Short-Term Variations and Long-Term Dynamics in Commodity Prices" Management Science 46, 893{911.

Singleton, Kenneth J. (2011) "Investor Flows and the 2008 Boom Bust in Oil Prices" Working paper, Stanford University.

Tang,Ke, and Wei Xiong (2011) "Index Investment and the Financialization of Commodities" Working paper, Princeton University.

Figure 1

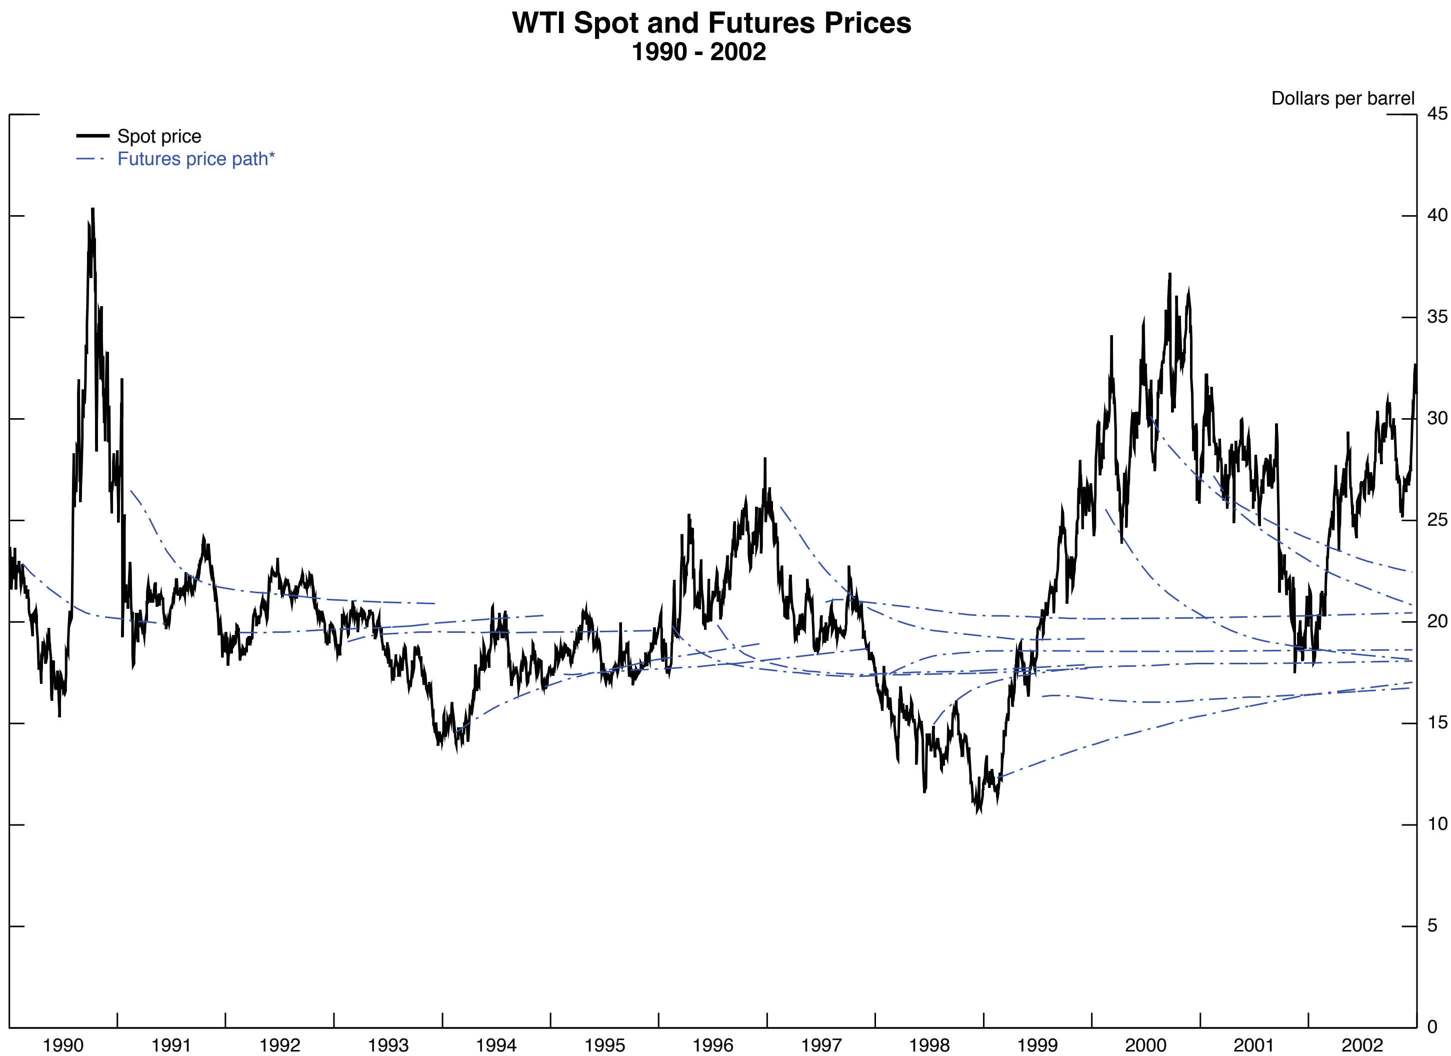

WTI Spot and Futures Prices
1990 - 2002
Dollars per barrel
Spot price
Futures price path*
45
40
35
30
25
20
15
10
5
0
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002

Figure 2

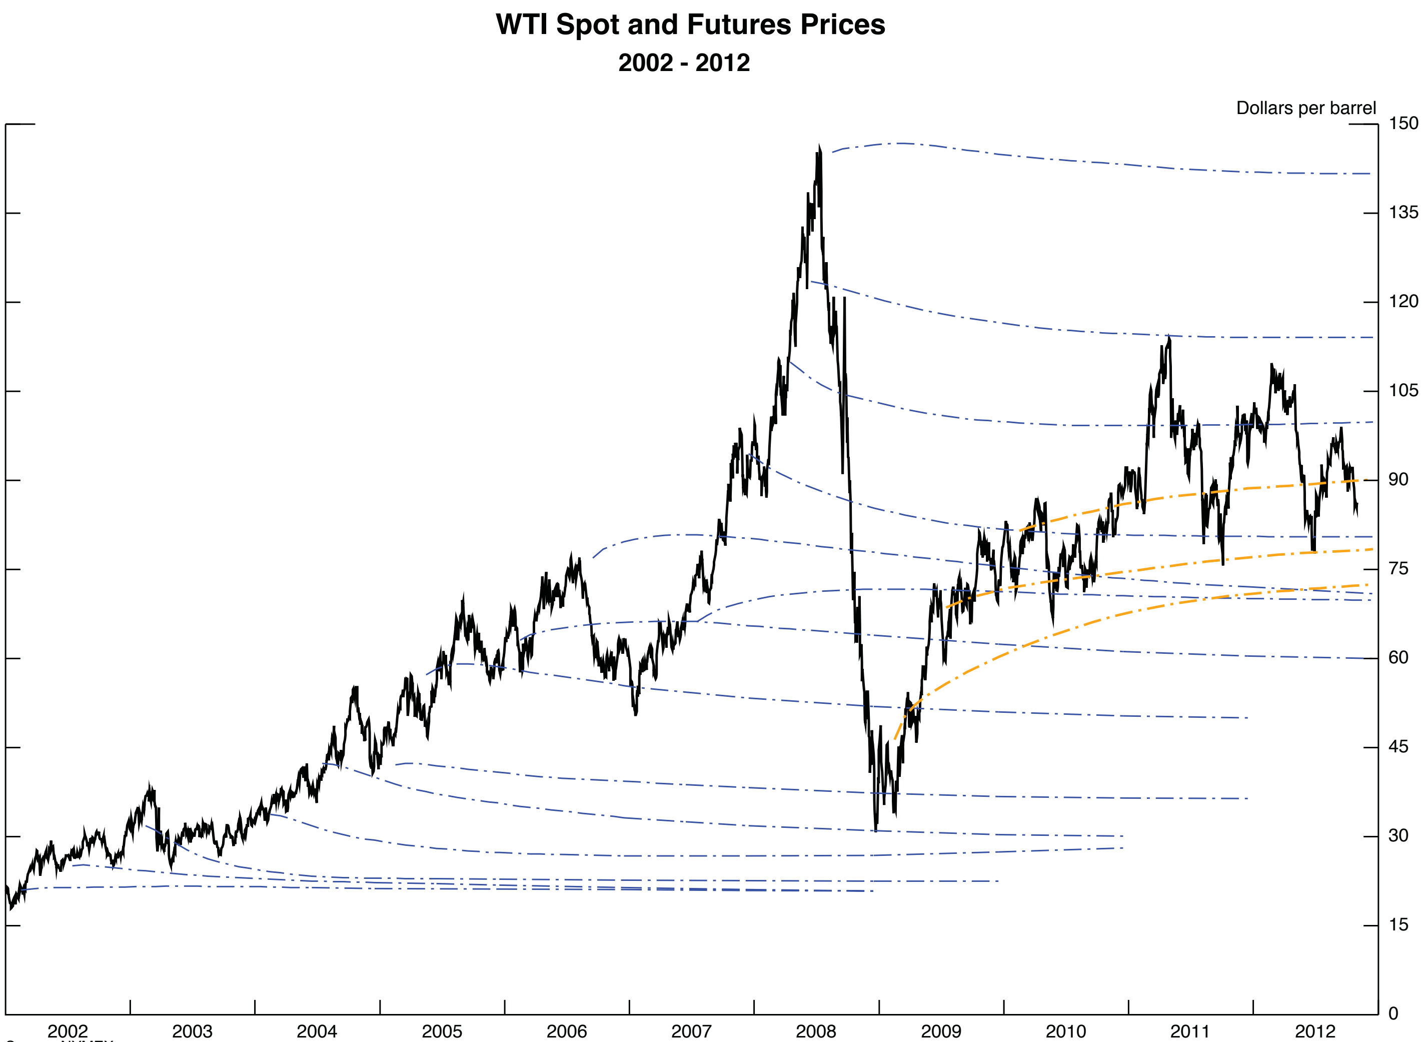

Source: NYMEX.
*Futures path as of first business day at 6-month intervals. The color of the futures paths is for clarity only.
Spot price is West Texas Intermediate crude oil, Cushing, OK.

Figure 3: Recursive Estimates of Structural Parameters

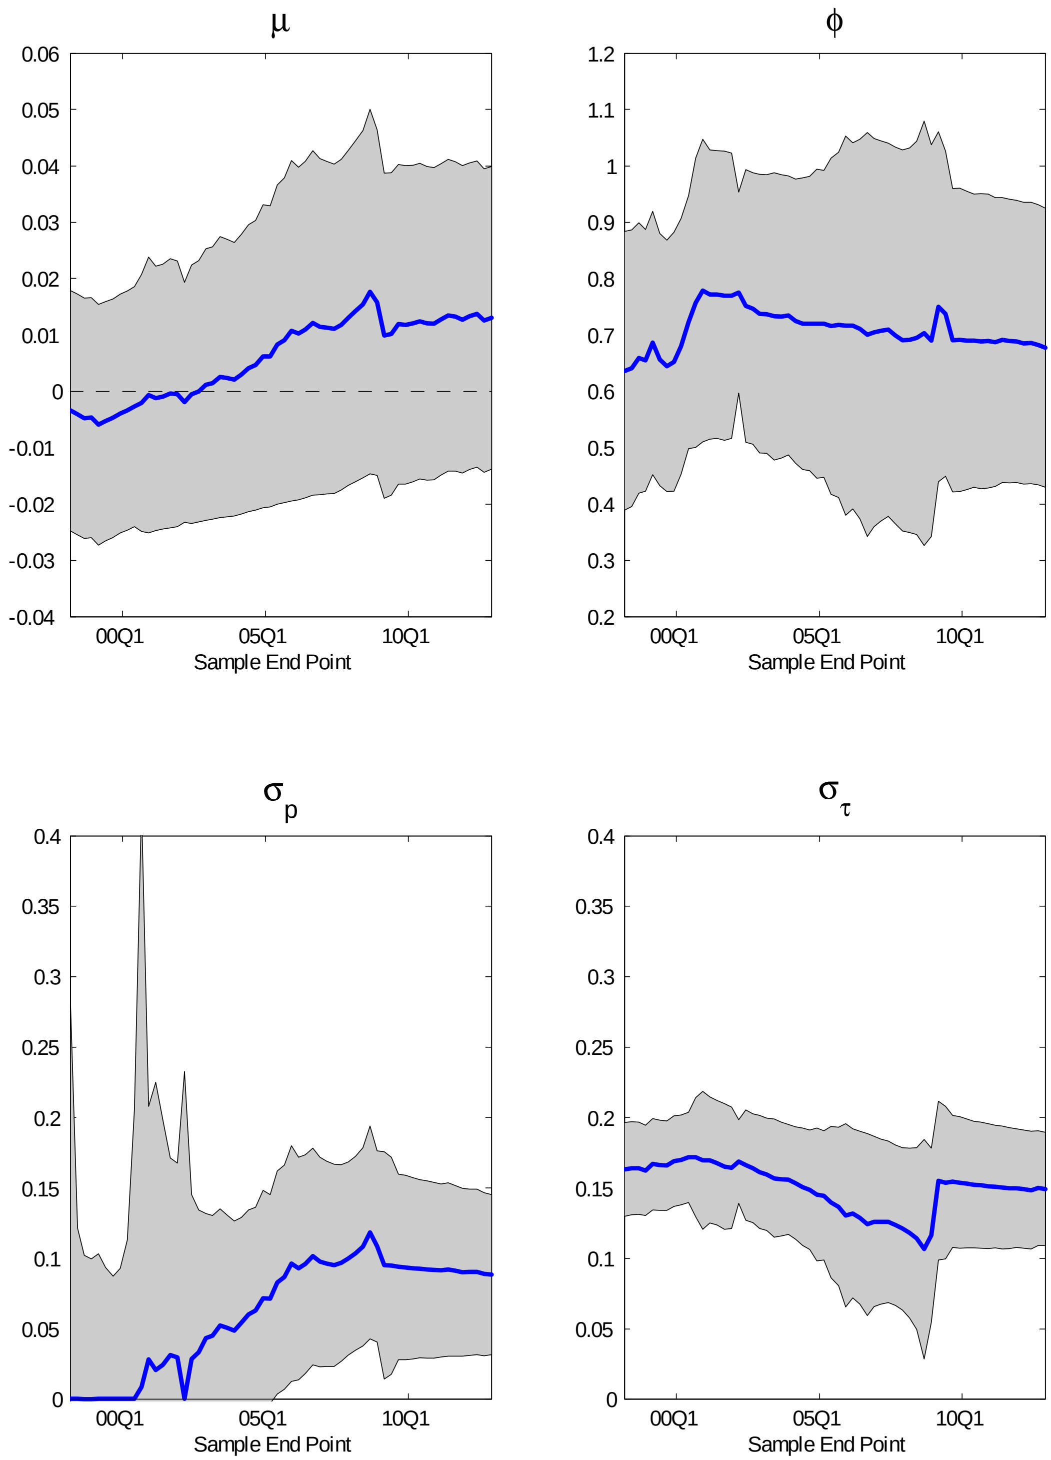

Figure 4: Futures Prices and Expected Futures Prices

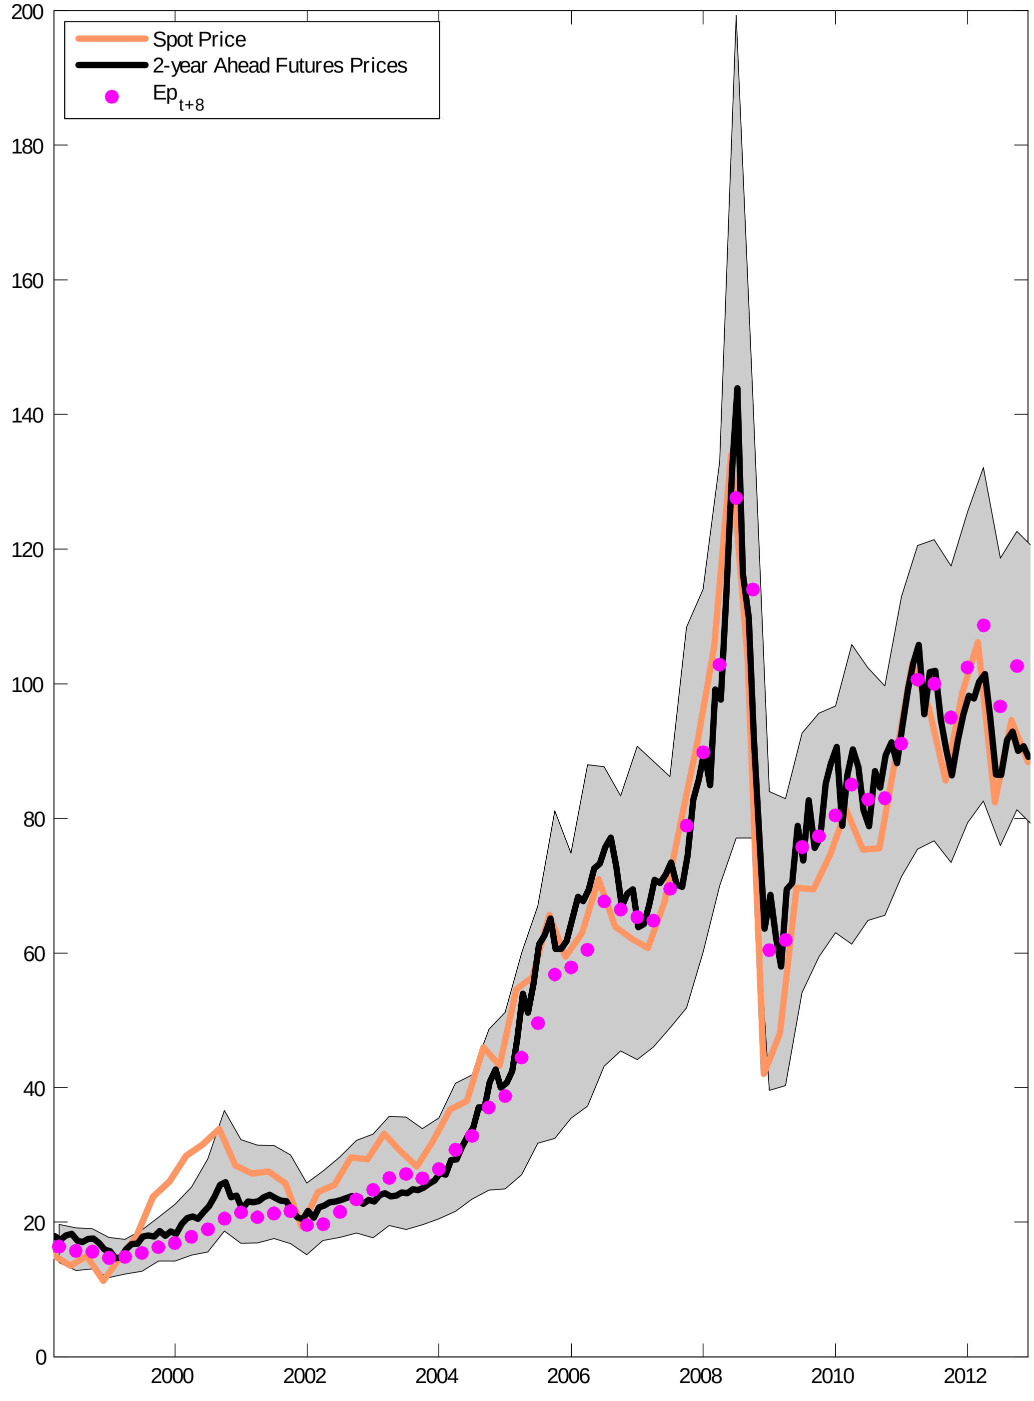

Figure 5: The Importance of time-varying Γ

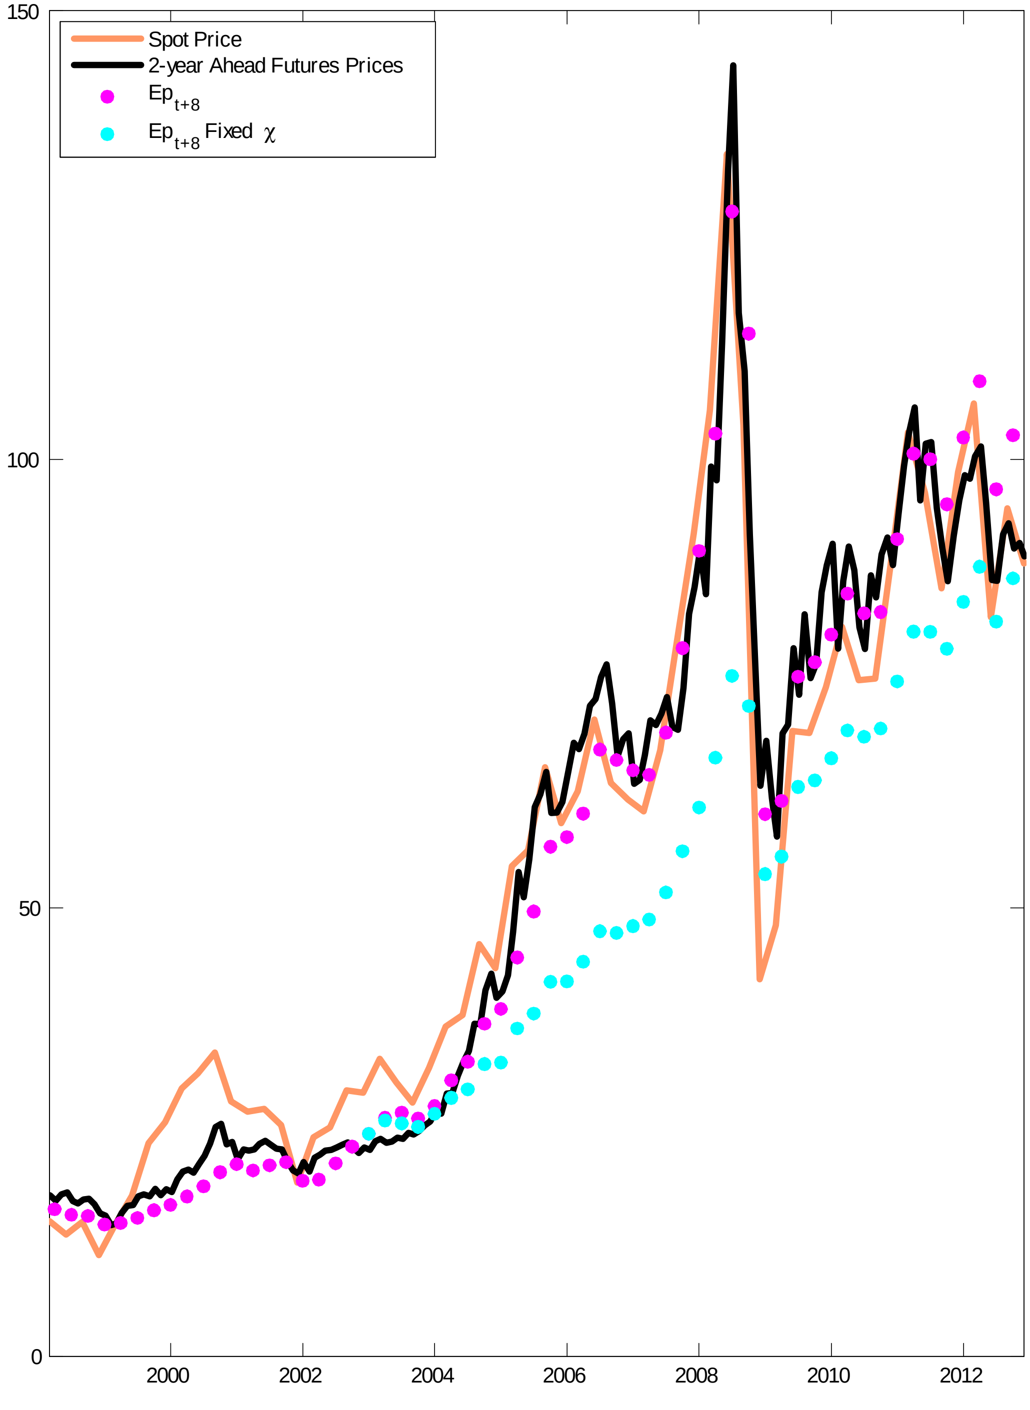

Figure 6: Expected Futures Prices for Constant Gain MLE

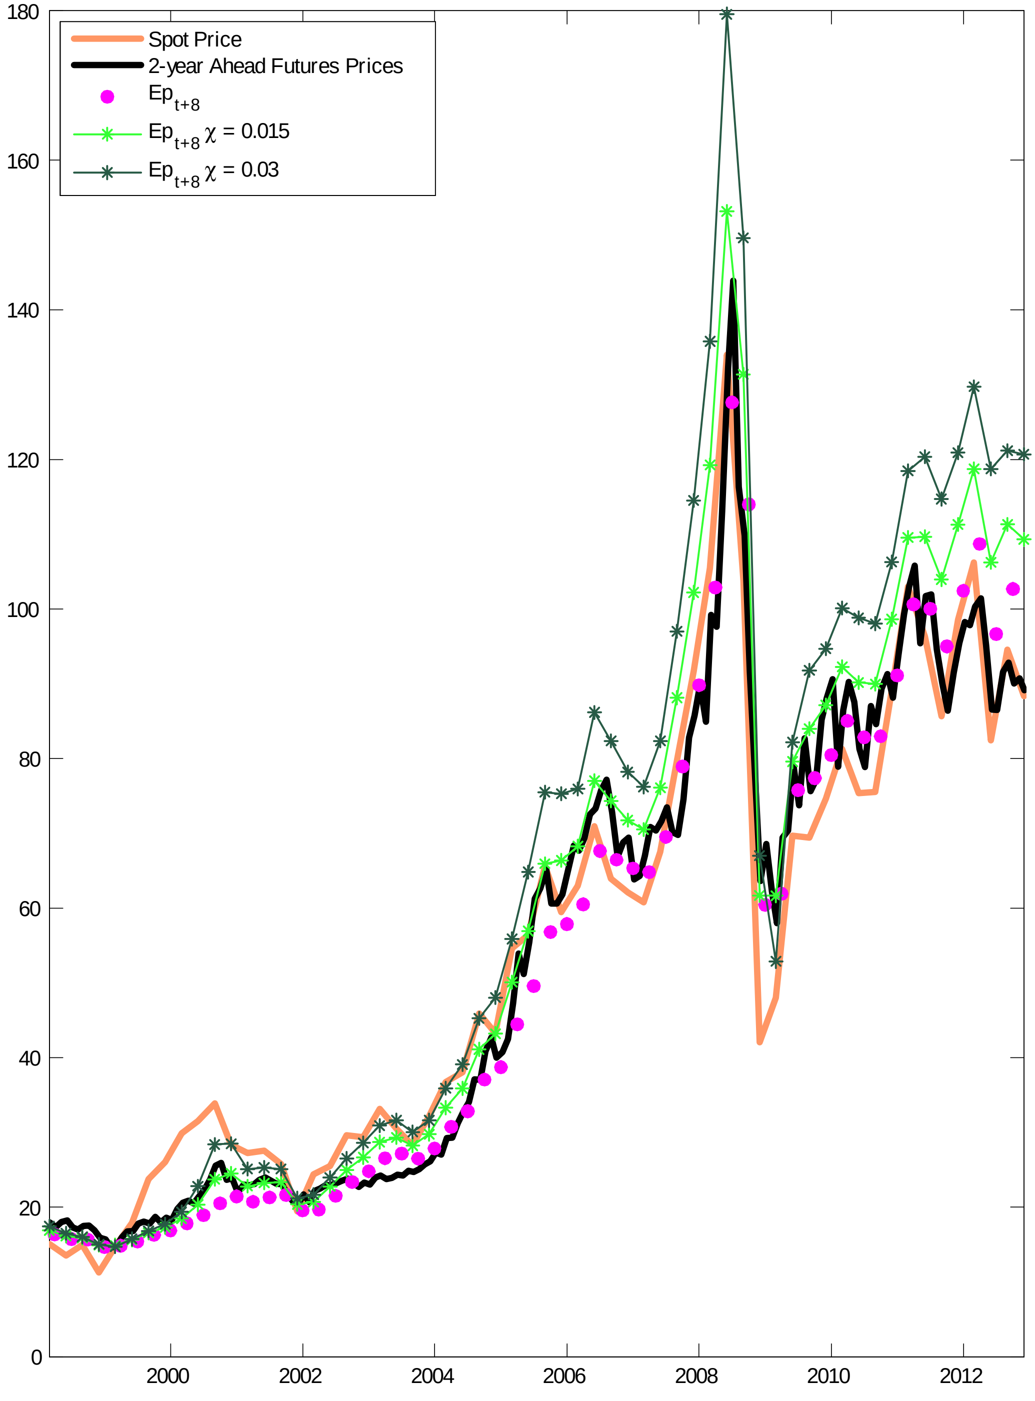

Figure 7: Coefficient Estimates for Constant Gain MLE

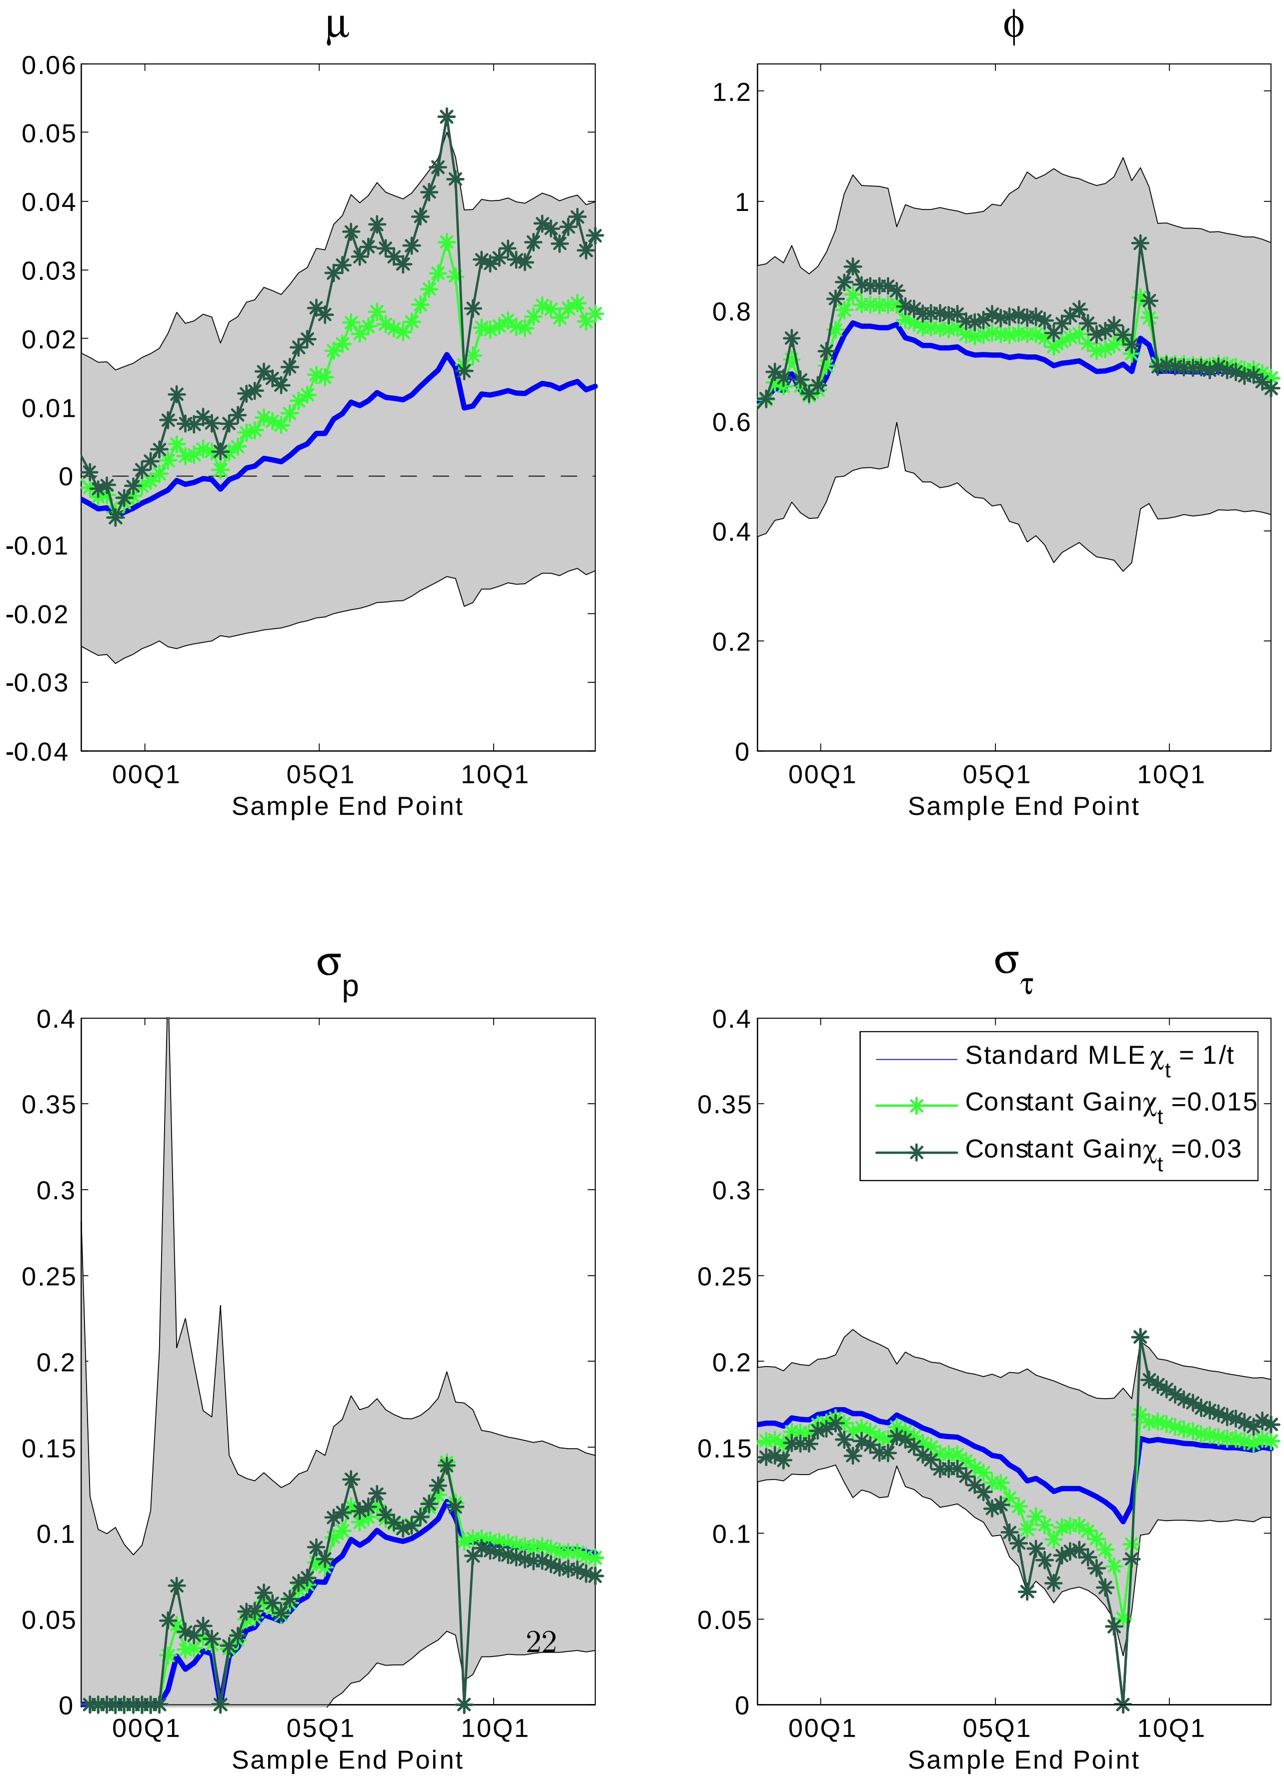

Figure 8: Variance Contribution to Δp_t For Constant Gain MLE

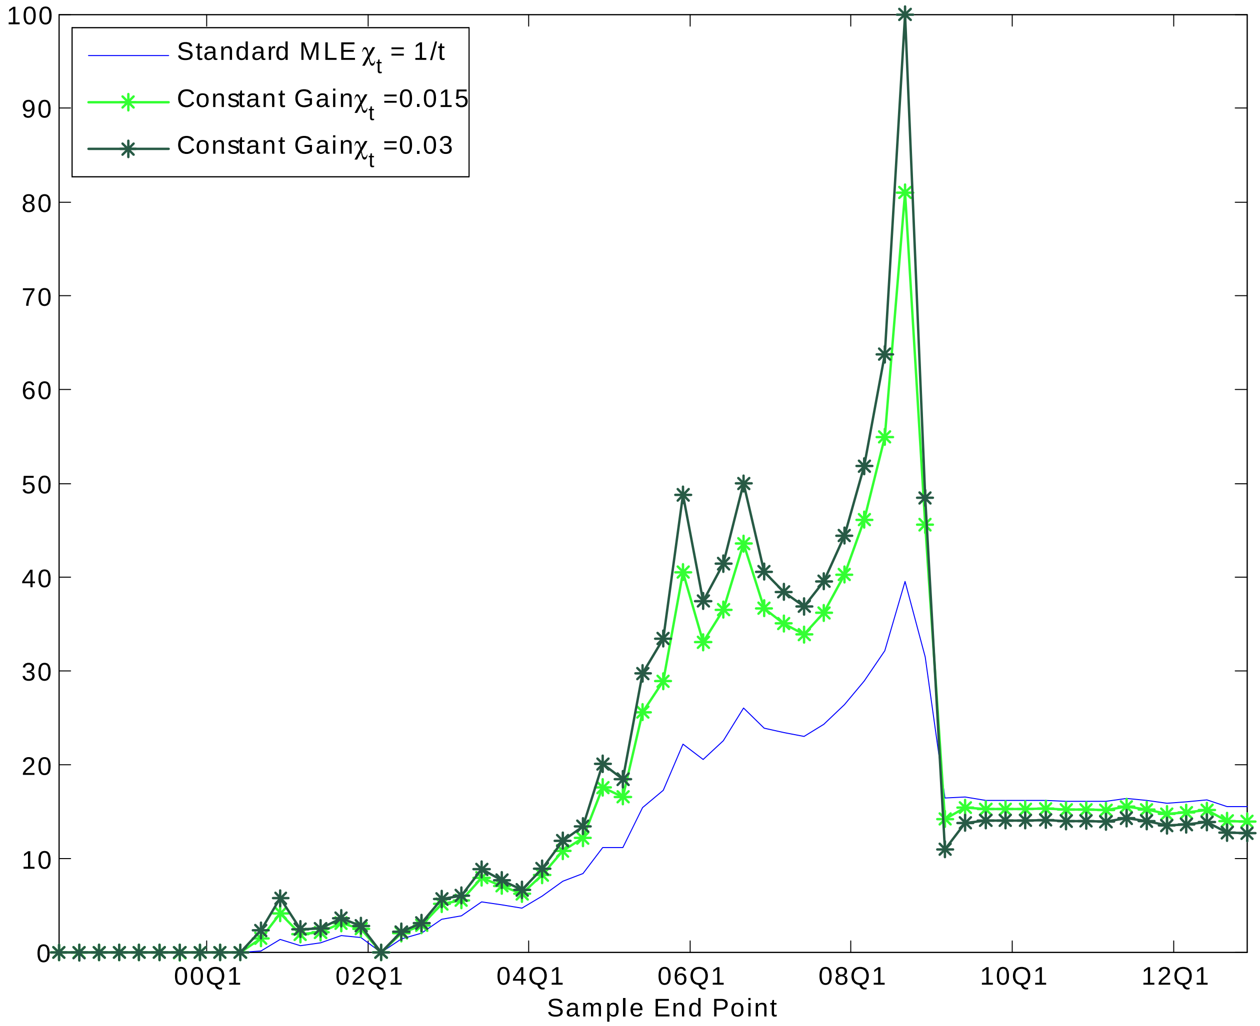

Figure 9: Recursive Estimates of β.